

Natural selection. II. Developmental variability and evolutionary rate

Steven A. Frank*

*Department of Ecology and Evolutionary Biology,
University of California, Irvine, CA 92697-2525 USA*

In classical evolutionary theory, genetic variation provides the source of heritable phenotypic variation on which natural selection acts. Against this classical view, several theories have emphasized that developmental variability and learning enhance nonheritable phenotypic variation, which in turn can accelerate evolutionary response. In this paper, I show how developmental variability alters evolutionary dynamics by smoothing the landscape that relates genotype to fitness. In a fitness landscape with multiple peaks and valleys, developmental variability can smooth the landscape to provide a directly increasing path of fitness to the highest peak. Developmental variability also allows initial survival of a genotype in response to novel or extreme environmental challenge, providing an opportunity for subsequent adaptation. This initial survival advantage arises from the way in which developmental variability smooths and broadens the fitness landscape. Ultimately, the synergism between developmental processes and genetic variation sets evolutionary rate^{ab}.

In evolutionary biology, environmentally induced modifications come under unfinished business ... There have been repeated assertions of both their importance and their triviality, a lot of discussion with no consensus. ... Yet the debate has continued over such concepts as genetic assimilation, the Baldwin effect, organic selection, morphoses, and somatic modifications. So much controversy over the span of a century suggests that a problem of major significance remains unsolved [1, p. 498].

I. INTRODUCTION

A single genotype produces different phenotypes. Developmental programs match the phenotype to different environments. Intrinsic developmental fluctuations spread the distribution of phenotypes. Extrinsic environmental fluctuations perturb developmental trajectory. These nonheritable types of phenotypic variation are common.

Nonheritable phenotypic variation is not transmitted through time. Thus, nonheritable variation would seem to be irrelevant for evolutionary change, which instead depends on the genetic component of variation. However, nonheritable phenotypic variation can, in principle, affect evolutionary rate. At first glance, that contribution of nonheritable phenotypic variation to evolutionary rate appears to be a paradox.

Many different theories, commentaries, and controversies turn on this paradox (Box 2). The literature has followed a consistent pattern. Detailed theories relate developmental variability to accelerated evolution. Counterarguments ensue. Listings of complicated examples

claim to support the theory. Refinements to the theory develop.

In the end, few compelling examples relate nonheritable phenotypic variability to evolutionary rate. The literature is hard to read. Enthusiasts extend the concepts and keep the problem alive. Through the enthusiasts' promotions, many have heard of the theory. But, in practice, few consider the role of nonheritable phenotypic variability in their own analyses of evolutionary rate. Almost everyone ignores the problem.

In this article, I emphasize simple theory that relates nonheritable phenotypic variability to evolutionary rate. Understanding the paradoxical relation between nonheritable phenotypic variability and evolutionary rate is an essential step in reasoning about many evolutionary problems.

This article is primarily a concise tutorial to the basic concepts (see Box 1). I briefly mention some of the history (Box 2) and recent, more advanced literature (Box 3).

II. SMOOTHING THE EVOLUTIONARY PATH

The distribution of phenotypes for a given genotype is called the *reaction norm*. All theories come down to the fact that a broad reaction norm smooths the path of increasing fitness. Once one grasps the smoothing process, many apparently different theories become easy to understand.

The next section gives the mathematical expression for the smoothing of fitness by the reaction norm. Fig. 1 explains the mathematics with a simple example.

A. The reaction norm smooths fitness

We need to track three quantities. First, fitness, $f(x)$, varies according to the particular phenotype expressed, x .

* email: safrank@uci.edu; homepage: <http://stevefrank.org>

^a doi: 10.1111/j.1420-9101.2011.02373.x in *J. Evol. Biol.*

^b Part of the Topics in Natural Selection series. See Box 1.

Box 1. Topics in the theory of natural selection

This article is part of a series on natural selection. Although the theory of natural selection is simple, it remains endlessly contentious and difficult to apply. My goal is to make more accessible the concepts that are so important, yet either mostly unknown or widely misunderstood. I write in a nontechnical style, showing the key equations and results rather than providing full derivations or discussions of mathematical problems. Boxes list technical issues and brief summaries of the literature.

Second, the phenotype expressed varies according to the reaction norm. Read $p(x|\bar{x})$ as the probability of expressing the phenotype x given a genotype with average phenotype \bar{x} .

Third, we must calculate $F(\bar{x})$, the expected fitness for a genotype with average phenotype \bar{x} . We obtain the expected fitness by summing up the probability, p , of expressing each phenotype multiplied by the fitness, f , of each phenotype. That sum is

$$F(\bar{x}) = \sum p(x|\bar{x})f(x), \quad (1)$$

taken over all the different phenotypes, x . We often measure x as a continuous variable. The sum is then equivalently written as

$$F(\bar{x}) = \int p(x|\bar{x})f(x)dx. \quad (2)$$

This equation shows how one averages the fitness, $f(x)$, for each phenotypic value, x , over the reaction norm, $p(x|\bar{x})$, to obtain the expected fitness of a genotype, $F(\bar{x})$. We label each genotype by its average phenotype, \bar{x} . The expected fitness of a genotype, $F(\bar{x})$, is what matters for evolutionary process [2].

The averaging of expected fitness over the reaction norm is the key to the entire subject. Averaging over the reaction norm, p , flattens and smooths the fitness function, f . This smoothing makes the curve for expected fitness, F , have lower peaks and shallower valleys than the original fitness curve, f . The smoothing of F changes evolutionary dynamics. The whole problem comes down to understanding how reaction norms smooth fitness, and the consequences of a smoother relation between genotype and fitness.

B. Example of continuous smoothing

Fig. 1 shows an example of smoothing with discrete distributions. It will often be convenient to consider smoothing of continuous variables. Fig. 2 shows an example. The following expressions describe the underlying mathematics.

Box 2. Historical overview

Schlichting and Pigliucci [3] and West-Eberhard [1] thoroughly review the subject. Here, I highlight a few key points in relation to this article. I treat learning and developmental plasticity as roughly the same with regard to potential consequences for evolutionary rate, although one could certainly choose to focus on meaningful distinctions.

In my own reading during the 1980s, I had found the relation between learning and evolutionary rate intriguing but confusing. Baldwin's [4] idea that learning can accelerate evolutionary rate seemed attractive. Mayr [5], in his monumental review of biological thought, also discussed various ways in which behavior or flexible developmental programs might alter evolutionary dynamics. Those ideas seemed potentially important, but it was not easy to grasp the essence. The literature at that time was not helpful, with a lot of jargon and sometimes almost mystical commentary mixed in with intriguing and creative ideas.

It was clear that learning could slow evolutionary rate. Different genotypes could, through learning, end with the same phenotype. Reducing the phenotypic distinction between different genotypes would generally slow evolutionary rate. The more intriguing problem concerns the origin of evolutionary novelty or the response to novel or extreme environmental challenge. Environmental novelty and acceleration of evolutionary response were the primary concern of Baldwin [4], Waddington [6, 7], and West-Eberhard [1]. My article also focuses on acceleration of evolutionary response.

Hinton and Nowlan [8] clarified the subject with their simple conclusion that:

Learning alters the shape of the search space in which evolution operates and thereby provides good evolutionary paths towards sets of co-adapted alleles. We demonstrate that this effect allows learning organisms to evolve *much* faster than their nonlearning equivalents, even though the characteristics acquired by the phenotype are not communicated to the genotype.

During the past few decades, the fundamental role of smoothed fitness surfaces in biology has not always been recognized as fully as it should be, in spite of several fine papers along that line (see Box 3). Interestingly, certain computer optimization algorithms take advantage of the increased search speed provided by a process similar to smoothed fitness landscapes [9, 10].

In Fig. 2, the reaction norm follows a normal distribution. In symbols, we write

$$p(x|\bar{x}) \sim \mathcal{N}(\bar{x}, \gamma^2),$$

which we read as the probability, p , of a phenotype, x , for a reaction norm centered at \bar{x} , follows a normal distribution with mean \bar{x} and variance γ^2 .

For fitness, we write in symbols

$$f(x) \sim \mathcal{N}(0, \sigma^2),$$

which we read as the fitness, f , of a phenotype, x , has the shape of a normal distribution with mean 0 and variance

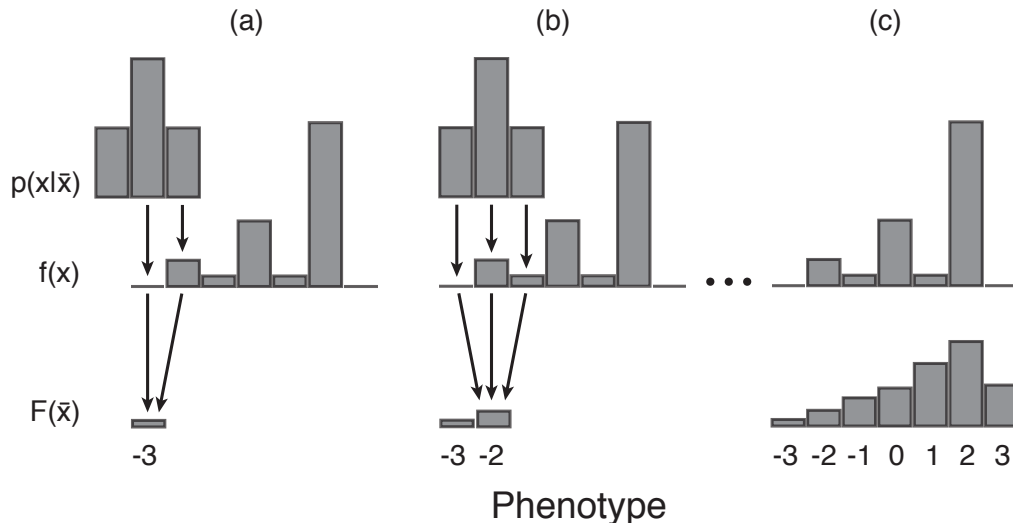


FIG. 1. The reaction norm smooths the fitness landscape. This simple example illustrates the calculation of the expected fitness for each genotype, following Eq. (1). (a) The calculation of expected fitness, $F(\bar{x})$, for the smallest average phenotype, $\bar{x} = -3$. For that average phenotype, the reaction norm, $p(x|\bar{x})$, shows the probabilities of expressing different phenotypes, x . In this case, the peak of the reaction norm matches the average value, and each phenotype ± 1 occurs half as often as the peak value. To get the expected fitness for a reaction norm centered at $\bar{x} = -3$, one sums up the probability $p(x|\bar{x})$ for each phenotype, x , multiplied by the fitness for each phenotype, $f(x)$. The arrows illustrate the summation. (b) The expected fitness, $F(\bar{x})$, for each increase in \bar{x} , is calculated by the same summation process, shifting the reaction norm to the right by one to get the proper value for each \bar{x} . (c) The full transformation is shown between the fitness for each phenotypic value, $f(x)$, and the expected fitness, $F(\bar{x})$, for each genotype with reaction norm $p(x|\bar{x})$ and average phenotype \bar{x} . The reaction norm smooths the multi-peaked fitness function, $f(x)$, into the single-peaked fitness function $F(\bar{x})$. Evolutionary dynamics depend on genotypic fitnesses, F . Thus, the reaction norm transforms fitness into a smooth function that allows a direct increasing path to the fitness peak from any starting value for average phenotype.

σ^2 . In this case, we assume the center of the fitness distribution is at a phenotypic value of zero to give a fixed point for comparison—any value to center fitness could be used. The important issue is that fitness falls off from its peak by the pattern of a normal distribution. The width of the fitness function is set by the variance parameter, σ^2 .

We can now use Eq. (2) to calculate the expected fitness of a genotype with average phenotype \bar{x} , yielding

$$F(\bar{x}) \sim \mathcal{N}(0, \gamma^2 + \sigma^2). \quad (3)$$

This equation shows that smoothing by the reaction norm, p , flattens and widens the shape of the fitness function by increasing the variance of the expression for F .

C. Evolutionary response to novel or extreme challenge

If a genotype expresses an average phenotype close to the maximum fitness, then a narrow reaction norm has higher fitness than a broad reaction norm. The lower plots of $F(\bar{x})$ in Fig. 2 illustrate contrasting widths of reaction norms. Near the peak, the average phenotype closely matches the optimum, and the narrower reaction norm has higher fitness. This advantage occurs because a

narrow reaction norm expresses fewer phenotypes in the tails, away from the optimum.

For genotypes with an average phenotype far from the maximum fitness, a broad reaction norm has higher fitness than a narrow reaction norm. Fig. 3 illustrates this advantage for broad reaction norms. In that figure, both reaction norms are centered at \bar{x} . Only those phenotypes above the fitness truncation point survive. The broad reaction norm produces some individuals with phenotypes above the truncation point, whereas the narrow reaction norm has zero fitness.

If the environment poses a novel or extreme challenge, the broad reaction norm wins. By contrast, in a stable environment for which the current average phenotype is close to the fitness optimum, the narrow reaction norm wins. Thus, extreme or novel environmental challenges or intense competition favor a broad reaction norm.

Haldane [11] made a similar point when he said: “Intense competition favors variable response to the environment rather than high average response. Were this not so, I expect that the world would be much duller than is actually the case.” Holland’s [12] emphasis on exploration versus exploitation is perhaps closer to the problem here. Broad reaction norms are favored when exploration of novel challenges dominates, whereas narrow reaction norms are favored when exploitation dominates. Fluctuating environments may also favor a broad




FIG. 2. Reaction norms and fitness for continuous phenotypes. Each column shows how the reaction norm, $p(x|\bar{x})$, smooths the fitness function, $f(x)$, to give the expected fitness, $F(\bar{x})$, for a genotype with average phenotype \bar{x} . The smoothing follows Eq. (2). These examples use normal distributions that lead to Eq. (3). (a) The solid and dashed reaction norms follow $\mathcal{N}(\bar{x}, 1/2)$ and $\mathcal{N}(\bar{x}, 5)$, respectively. Fitness, $f(x)$, has the shape of a normal distribution with vanishingly small variance, $\mathcal{N}(0, \sigma^2 \rightarrow 0)$. Thus, expected fitness, $F(\bar{x})$, is the same as the reaction norm. (b) The same structure as in (a), except that $f(x)$ is much wider, following $\mathcal{N}(0, 7)$. Thus, $F(\bar{x})$ now has curves $\mathcal{N}(0, 7.5)$ and $\mathcal{N}(0, 12)$ for solid and dashed curves, respectively. In each plot, the baseline is set to 4.3% of the peak in that plot. The baseline truncates phenotypes with low vigor, setting their fitnesses to zero.

reaction norm to increase the chance of matching whatever is favored at any time [2]. Here, I focus on constant challenges to extreme or novel environments.

D. Smoothly increasing fitness path in a multipeak fitness landscape

Much discussion in evolutionary theory concerns how populations shift from a lower fitness peak to a higher fitness peak [13]. For example, in the fitness landscape, $f(x)$, of Fig. 4b, a population starting on a lower peak must evolve through a valley of lower fitness in order to follow an increasing path to a higher fitness peak. Natural selection typically follows a path of increasing fitness, so a population may be trapped on a lower peak.

Most evolutionary analyses use a fitness landscape that relates phenotype, x , to fitness, $f(x)$. However, the proper measure should relate the average phenotype of a genotype, \bar{x} , to the expected fitness, $F(\bar{x})$ [2].

A sufficiently broad reaction norm smooths a multi-peak fitness landscape, $f(x)$, into a smooth landscape,




FIG. 3. Novel environmental challenge or intense competition favors a broad reaction norm. In this example, both the broad and narrow reaction norms are centered at \bar{x} . Phenotypes above the truncation point survive. Phenotypes below the truncation point die. None of the phenotypes for the narrow reaction norm are above the truncation point, so all die. Some of the phenotypes of the broad reaction norm survive. Those surviving phenotypes may evolve so that their average phenotype, \bar{x} , moves toward the truncation point, improving fitness over time. Improvement occurs if there is genetic variation for the average phenotype, \bar{x} , of the reaction norm.

$F(\bar{x})$, with a single peak (Fig. 4c). A broad reaction norm will typically perform badly near a fitness peak, but allow much more rapid evolutionary advance to a higher fitness peak. Once again, we see that broad reaction norms exploit current fitness opportunities relatively poorly but gain by enhanced exploration and achievement of novel adaptations.

III. DIMENSIONALITY AND DISCOVERY

The reaction norm may be generated randomly by perturbations in development. If so, then exploration of the fitness landscape by a broad reaction norm is a type of random search. Figs. 3 and 4 show that random search can greatly increase the rate of adaptation, particularly to novel environmental challenges.

Those previous examples showed the reaction norm and fitness both varying across a single dimension. A broad reaction norm spreads phenotypes along that single dimension, increasing the chance that some individuals will have high fitness.

Now consider the much more difficult search problem that arises in higher dimensions [14]. Suppose, for example, that adapting to a novel environmental challenge requires multiple phenotypic changes to work together in a harmonious way. Think of each particular phenotypic change as a trait in its own dimension, so that the search now occurs in multiple dimensions. If the reaction norm simply generates random phenotypes in each dimension, then there is little chance of getting simultaneous match-

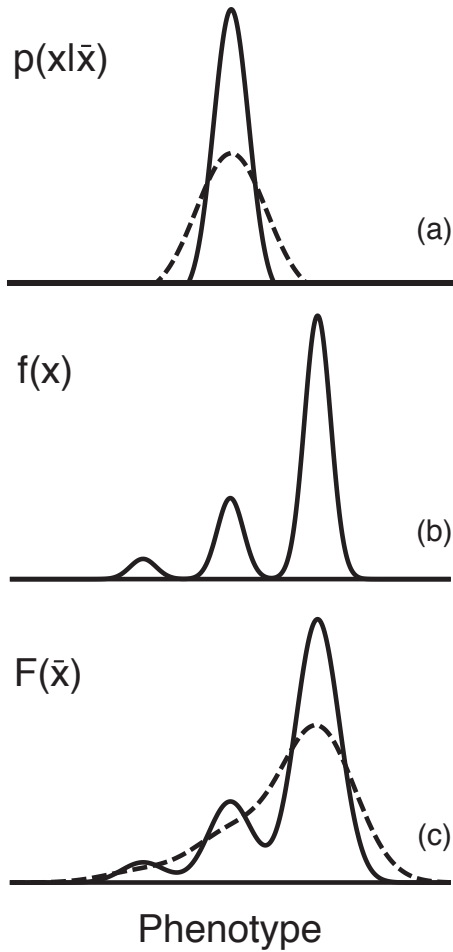


FIG. 4. A broad reaction norm smooths a multipeak fitness landscape. (a) The dashed curve shows the broader reaction norm, $p(x|\bar{x})$. (b) The fitness landscape for each particular phenotype, $f(x)$, has multiple peaks. (c) The broad reaction norm smooths the fitness landscape to a single peak for the relation between the average phenotype for a genotype, \bar{x} , and fitness, $F(\bar{x})$. In this example, the narrow and broad reaction norms follow $\mathcal{N}(0, \gamma^2)$ distributions with variances of 0.04 and 0.16, respectively. Fitness is given by $f(x) = \sum_{i=-1}^1 (3|1+i|^2 + 1)\mathcal{N}(i, \sigma^2)$, with $\sigma^2 = 0.0225$. The value of $F(\bar{x})$ is calculated from Eq. (2), yielding the expression for $f(x)$ in the prior sentence with the variance replaced by $\sigma^2 + \gamma^2$. The baseline truncates small values.

ing phenotypes in multiple dimensions.

To visualize the multidimensional problem, begin with the one-dimensional fitness landscape in Fig. 4b. Now consider two phenotypic dimensions. Assume that fitness concentrates along one dimension, as in Fig. 5a. In that plot, only a narrow band of phenotypes along the second phenotypic dimension produces viable individuals. In the first dimension, fitness rises and falls along the same peaks and valleys as in Fig. 4b. Thus, both figures show essentially the same fitness landscape, but in the second case the nearly one-dimensional landscape is embedded in a second dimension (fitnesses scale loga-

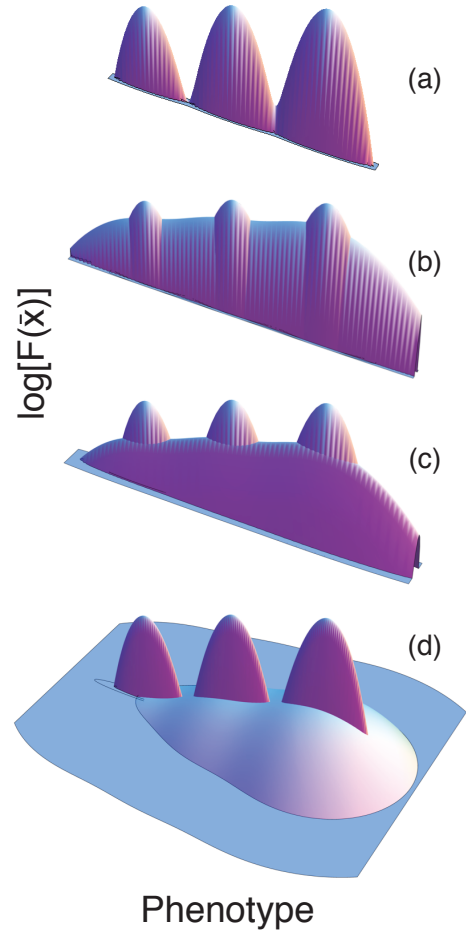


FIG. 5. A broad reaction norm performs poorly when fitness is concentrated in a lower dimension. (a) The bivariate analogy of the fitness landscape in Fig. 4b, scaled logarithmically. The primary dimension has variance $\sigma_1^2 = 0.0225$ corresponding to standard deviation $\sigma_1 = 0.15$, as in Fig. 4b. The secondary (narrow) dimension has standard deviation $\sigma_2 = 0.1\sigma_1$. (b) Fitness landscape smoothed by a reaction norm concentrated in the same dimension as fitness. The variance of the reaction norm in the primary dimension is $\gamma_1^2 = 0.16$, and standard deviation is $\gamma_1 = 0.4$, as in the dashed reaction norm of Fig. 4a. The standard deviation in the secondary dimension is $\gamma_2 = 0.01\gamma_1$. The smoothed fitness surface rises steadily to a peak along its ridge in the primary dimension, tracing the same path as the dashed curve in Fig. 4c. (c and d) Increasingly broad reaction norms in the secondary dimension with standard deviations of $0.1\gamma_1$ and γ_1 , respectively. The baseline truncates small fitness values, which are considered inviable.

arithmically in Fig. 5).

In two dimensions, the reaction norm will smooth phenotypes along both trait axes. When the reaction norm varies mostly along the same dimension as the variation in fitness, as in Fig. 5b, then we obtain the same smoothing as in one dimension (dashed curve of Fig. 4c). When the reaction norm varies in both directions, as in Fig. 5d, then the smoothed surface has very low fitness even at its




FIG. 6. Decline in fitness with an increasingly broad reaction norm away from the primary concentration of fitness. The dashed lines show the fitness associated with the high, medium and low peaks of the fitness landscape in Fig. 5a. The solid curve shows the highest point of the fitness functions smoothed by the reaction norms of Fig. 5b–d, with standard deviation of the reaction norm increasing in the secondary dimension. The secondary dimension standard deviation value shown in the plot gives the amount by which the primary dimension standard deviation is multiplied in the second dimension. The multipliers 0.01, 0.1, 1.0 correspond to the three smoothed fitness surfaces in Fig. 5b–d. When the secondary dimension is narrow, for example, reduced in width by a factor 0.01, then the smoothed fitness peak is higher than the intermediate fitness peak of the unsmoothed landscape, as in Fig. 4c. As width in the secondary dimension increases, the cost of exploring in a dimension away from the concentration of fitness causes the peak of the smoothed fitness landscape to drop very low, illustrating the very high cost of exploring in more dimensions than the concentration of fitness.

peak. The low fitness occurs because the randomly generated reaction norm produces phenotypes spread across two dimensions. Most of those phenotypes fall off of the one dimensional concentration of fit phenotypes.

In general, when the dimensionality of the reaction norm exceeds the dimensionality of the fitness concentration, then a random search process is inefficient. The cost of exploration is so high that even the best average phenotype for a genotype has fitness, $F(\bar{x})$, lower than the lowest peak of the fitness landscape, $f(x)$. Here, \bar{x} and x represent multidimensional phenotypes. Fig. 6 illustrates the cost of exploration in relation to the spread across dimensions.

In summary, if the space of possible trait combinations spreads over greater dimensions than the concentration of fitness, then randomly generated variations will produce mostly worthless variants. The search cost is high, and average performance for a widely spread reaction norm is low. The smoothed fitness surface may have a steadily rising path to its fitness maximum from many initial points, but the height of the fitness peak is so low that a broad reaction norm will often be strongly selected against.

The following sections describe two processes that may offset the high cost of developmental variation. First, the broad search space may be covered by genetic variants,

Box 3. Recent literature

Ancel [15] analyzed smoothed fitness surfaces and the consequences for evolutionary rate. She emphasized three important points.

First, learning accelerates evolution only under certain conditions. The examples in the text illustrate this point by showing that learning mainly accelerates evolution through discovery of viable phenotypes or in the smoothing of a multi-peaked fitness surface. Otherwise, the smoothing of fitness surfaces may lower the maximum fitness that can be attained, reducing the slope and the evolutionary rate to the peak.

Second, although learning may accelerate evolution, it is not necessarily true that learning evolved because it accelerates evolution. The evolutionary consequence of a trait is distinct from whether or not the trait evolved because of its potential to alter subsequent evolutionary dynamics. The literature discusses this distinction under the topic of evolvability. Evolvability has developed into a large subject of its own [16–20]. Holland’s [12] distinction between exploration and exploitation captures aspects of the later developments on evolvability.

Third, Ancel noted historical precedents for the idea that phenotypic variance may eliminate otherwise uncrossable valleys in fitness landscapes [21–23].

A large literature develops issues related to Ancel’s three points and the broader problems of how reaction norms affect fitness surfaces. I list a small sample [14, 24–42].

West-Eberhard [1] discusses many empirical issues and examples. Recent studies in evolutionary biology provide new data or summaries of the literature [43, 44]. Articles in microbiology and cancer research have also developed the relation between nonheritable phenotypic variation and evolutionary process [45–53].

In the text, I discuss the synergism between genetic and developmental variation. I am not aware of literature related to that issue. However, given the many papers on the general topic, the synergism between genetics and development may have come up previously.

with developmental variation searching only the local regions around each genotypic variant.

Second, developmental variation may be biased in a way that tends to match the environment. If a developmental or learning process brings the phenotype close to the concentration of fitness in a high dimensional space, then some additional random variation can greatly increase the rate of adaptation. In this case, the fitness surface is smoothed to provide a steady path of increasing fitness, and the developmental bias that brings the center of the phenotypic distribution close to the fitness concentration mitigates the large cost of search in high dimensional phenotypic spaces.

IV. SYNERGISM BETWEEN PHENOTYPIC AND GENETIC VARIATION

A broad reaction norm may enhance survival and subsequent opportunity for improved fitness. But those ben-

efits arise only when a genotype is sufficiently close to a fitness peak. Fig. 7 illustrates the problem.

Fig. 7a shows the fitness peak in a novel environment. The dots show the locations of alternative genotypes, placed by their average phenotypes in two dimensions. Neither genotype has positive fitness. Both will die out. In that plot, the fitness peak is the direct fitness landscape, unsmoothed by a reaction norm. In Fig. 7b, the reaction norm is relatively narrow, smoothing the fitness landscape. But that smoothing is not enough to place either genotype on the nonzero fitness surface. Both genotypes still die out.

The broader reaction norm in Fig. 7c smooths the fitness surface more widely. That additional smoothing allows the nearby genotype to survive. Subsequent small genetic variations would allow natural selection to drive the surviving population up the path of increasing fitness to the fitness peak.

The distant (red) genotype cannot survive even with the broad reaction norm of Fig. 7c. The contrast between the nearby and distant genotypes emphasizes a key point. A genotype must be sufficiently close to the nonzero part of the smoothed fitness surface in order for the developmental variation of the reaction norm to allow survival—the touching of the fitness surface. If a genotype touches the fitness surface, then it can seed a population in which small genetic variations allow subsequent adaptation by climbing the surface to the peak.

In a high dimensional space, any single genotype is unlikely to be located sufficiently close to a fitness peak after a significant change in the environment or in response to an unpredictable challenge. Synergism between genetic variation and the phenotypic variation of reaction norms provides one solution to this search problem.

Fig. 8 illustrates synergism between genetic and phenotypic variation. The dots represent different genotypes. Each genotype has a different combination of average phenotypic values in two dimensions. The array of dots shows the genetic diversity in the population. The smoothed fitness surface has the same fitness peak and reaction norm shape as in Fig. 7c. In Fig. 8, the location of the fitness surface varies in the different plots, illustrating different environmental challenges. No matter where the newly favored fitness surface arises upon environmental challenge, the genetic diversity in the population provides at least one genotype on the nonzero part of the novel fitness surface. Those genotypes on the surface can survive the novel challenge. Subsequent small genetic variations around a surviving genotype allow the population to evolve up the fitness surface to the peak set by the novel environmental challenge.

Synergism between genetic and phenotypic variation divides the adaptive search problem into three parts. Genetic variation covers widely separated locations in the phenotype space. Reaction norms cover the phenotype space around each genotype. Any genotype on a nonzero part of a novel fitness surface can survive and subsequently adapt by small genetic variations and natural

Box 4. Vertebrate immunity

Invading pathogens present a vast diversity of foreign molecules that must be recognized. The vertebrate adaptive immune system develops antibodies by synergism between phenotypic and genetic variation, following the general three-part search process described in the text [54, 55].

First, to generate genetic diversity, B cell lineages within the body undergo programmed genetic recombination early in life. That recombination yields genetically distinct cellular clones. Each clone produces a distinct antibody.

Second, each antibody type from this initial diversity tends to bind relatively weakly to a variety of foreign antigens. In this regard, the original or “natural” antibodies trade the cost of weak binding for the benefit of a phenotypically diverse response—a broad reaction norm. Upon challenge with a foreign antigen, those B cells with matching antibodies are stimulated to expand clonally. That clonal expansion can be thought of as survival and reproduction of those genotypes that land on the fitness surface imposed by the unpredictable invader.

Third, the weakly binding antibodies undergo a programmed round of hypermutation to the antibody binding site and selection favoring variants that bind more tightly to the foreign antigen. This affinity maturation produces tightly binding and highly adapted antibodies in response to the novel challenge. Put another way, the initially stimulated antibodies on the edge of the “fitness surface” climb the surface toward the fitness peak.

In the process of climbing the fitness peak by local genetic variation and natural selection, the refined antibodies match more closely to the environmental challenge. In particular, the refined antibodies narrow their reaction norm by increasing their binding affinity for close matches and reducing their binding affinity for slightly mismatched binding.

In summary, the ability of the adaptive immune system to respond to the huge diversity of potential challenges depends on its synergism between genetic variability and the reaction norm. The initial natural antibodies arise from genetically diverse clones produced by recombination. That genetic diversity by itself could not cover the huge space of possible challenges. The broad reaction norm around each genetic variant allows protection against novel challenge. Once partial recognition is achieved through the natural antibodies, the system refines the match locally by affinity maturation.

selection. See Box 4 for an example of the synergism between genetic and nonheritable phenotypic variation.

V. MATCHING THE ENVIRONMENT BY PLASTICITY OR LEARNING

To survive a novel environmental challenge, a phenotype must be near the nonzero part of the new fitness landscape. A population may survive by having a variety of genotypes that produce different phenotypes, increasing the chance that at least one of the phenotypes will be close to a new fitness peak. Alternatively, a single genotype may be able to produce diverse phenotypes

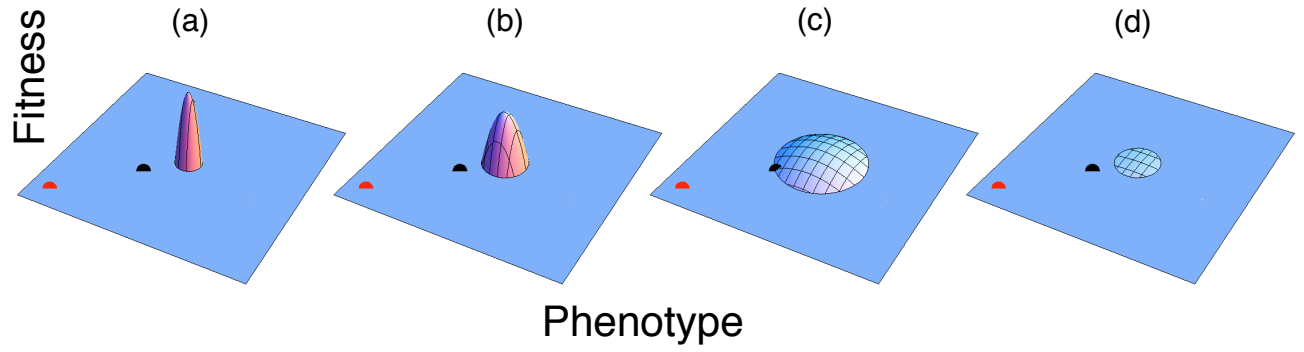


FIG. 7. Fitness in two phenotypic dimensions after challenge by a novel or extreme environment. The peak corresponds to the favored phenotype after environmental challenge. The black dot shows the average phenotype for a genotype near the new fitness peak. The red dot shows the average phenotype for a genotype relatively far from the new peak. A sufficiently broad reaction norm allows the nearby genotype to survive, providing an opportunity for natural selection to drive the population up the smoothly increasing path to the new fitness peak. By contrast, the distant genotype cannot survive the environmental challenge. (a) The fitness landscape, $f(x)$, showing the direct relation between phenotype and fitness when not smoothed by a reaction norm. (b) A relatively narrow reaction norm smooths the fitness peak, $F(\bar{x})$, but not sufficiently to allow the nearby genotype to survive. (c) A broader reaction norm allows the nearby genotype to survive, with subsequent opportunity for natural selection to drive the population to the peak. (d) An increasingly broad reaction norm causes the smoothed fitness peak to sink mostly below the fitness truncation level, so that the nearby genotype cannot survive. All plots show a bivariate normal fitness surface with mean $(1/2, 1/2)$ and variance $\sigma^2 + \gamma^2$, with $\sigma^2 = 0.01$ for the fitness landscape, and γ^2 for the reaction norm of 0, 0.01, 0.07, 0.11 for plots left to right. The heights are the natural logarithm of fitness, with a truncation base of $\log(10)$. The nearby black dot is at $(17/40, 17/40)$, and the far red dot is at $(11/40, 11/40)$.

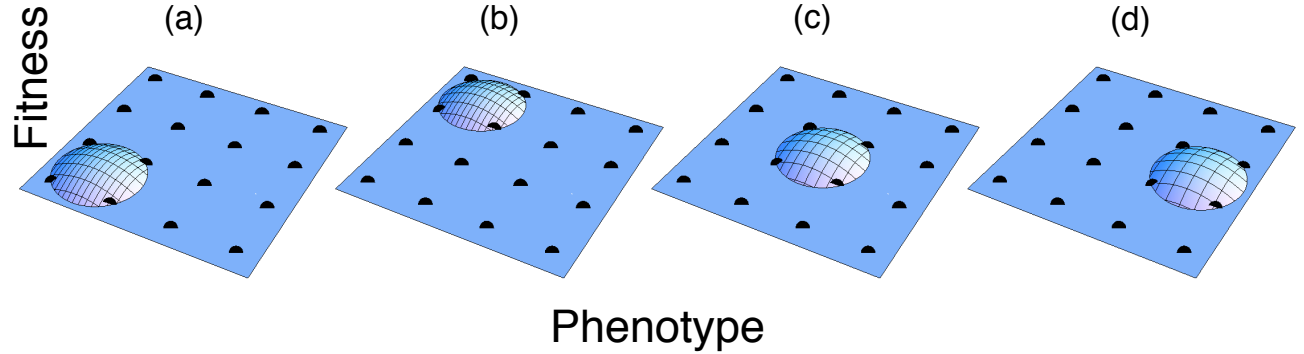


FIG. 8. Synergism between the reaction norm and genetic variation allows rapid adaptation to novel or extreme environments. In this case, a population has multiple genotypes, each genotype located at one of the dots. The smoothed fitness surface is sufficiently broad that, for any location of the fitness peak after environmental challenge, the fitness surface touches at least one of the genotypes. The genotypes that touch the fitness surface survive, allowing the surviving population the potential subsequently to evolve up the fitness surface to the new peak. The dots are located at all bivariate pairs from $\{(5 + 6i)/40, (5 + 6j)/40\}$ for $i, j = 1, \dots, 4$. The fitness surface has variance $\sigma^2 + \gamma^2 = 0.08$, as in Fig. 7c.

by matching phenotypic expression to the particular environment. The developmental flexibility to match environments may arise by phenotypic plasticity or learning.

Plasticity or learning may not be able to match exactly a novel or extreme environmental challenge. But if a developmental response to the environment can move the phenotype sufficiently close to the nonzero part of the new fitness landscape, then the genotype may survive and subsequently adapt [1, 4, 6]. Developmental flexibility is simply another process that alters the shape of the fitness surface.

The adaptive search problem has three phases, similar to the three aspects of search described in the prior

section. First, partially matching expression to the environment brings the phenotype close to the new fitness landscape. Second, random perturbations of phenotype occur around the location set by the process of environmental matching. Third, any genotype on a novel fitness surface can survive and subsequently adapt by small genetic variations and natural selection.

Fig. 9 illustrates the three aspects of adaptive search. Suppose a genotype expresses phenotypes centered at \bar{x} . In the first aspect of adaptive search, a genotype can adjust phenotypic expression to match the environment. The possible range of phenotypes varies from $\bar{x} - c$ to $\bar{x} + c$. The phenotype expressed by environmental matching is




FIG. 9. Plasticity or learning provides a partial match of phenotype to novel or extreme environmental challenge. A genotype's default expression has average phenotypic value \bar{x} . An individual can modify average trait expression in response to the environment. The average expressed phenotype can be any value in the range $\bar{x} \pm c$. The environmental challenge defines the fitness landscape that relates phenotype to fitness, here shown as the solid peak. Phenotypic expression, modulated by a match to the environment, shifts the phenotype to $\bar{x} + c$. However, the fitness associated with a phenotype of $\bar{x} + c$ is zero, because that value remains outside the range of viable phenotypes. Suppose $\bar{x} + c$ is the average phenotype expressed, and random perturbations of expression cause variability in phenotype around that average value. The random component of phenotype smooths the fitness landscape, leading to the dashed fitness surface. The expressed average phenotype, $\bar{x} + c$, now falls within the smoothed fitness surface, allowing the genotype to survive. Subsequent adaptation may allow improved fitness, by altering the range of phenotypes that can be expressed so that a match to the fitness peak may be achieved.

the new average value, around which random perturbations may occur. In the figure, the solid peak shows the fitness landscape imposed by a novel or extreme environmental challenge. The example genotype can come close to the new peak by modulating expression to produce an average phenotype of $\bar{x} + c$. However, if no random variation occurs around $\bar{x} + c$, that phenotype falls outside the range of phenotypic values that can survive.

In the second aspect of adaptive search, the genotype may produce phenotypes randomly distributed around the mean value of $\bar{x} + c$. Those random fluctuations smooth the fitness landscape, shown by the dashed curve. The average phenotype $\bar{x} + c$ can now survive. Matching the environment allowed expression of mean phenotype $\bar{x} + c$, and random fluctuations in phenotype smoothed the nearby landscape sufficiently.

Once the genotype achieves survival, the third phase of adaptation may proceed. In this case, the mean phenotype $\bar{x} + c$ has low fitness on the dashed fitness surface. But the fitness surface has a smoothly increasing path to the peak of maximum fitness. Genetic variations in the genotype may shift the range of phenotypes that can be produced, allowing natural selection to drive the population up the fitness surface to the peak.

Box 5. A simple prediction

Frequent exposure to novel or extreme environmental challenge favors a greater capacity for evolutionary response. Faster evolutionary response may be achieved by enhanced generation of genetic variability. For example, in experimental evolution studies, strong competition favors strains of bacteria with a high mutation rate [56]. Strong competition can be thought of as a form of extreme challenge. However, interference between genetically distinct competing clones complicates the association between mutation rate and evolutionary rate [57].

Faster evolutionary response may also be achieved by enhanced phenotypic variability through learning or plasticity. However, it may be difficult to test empirically whether learning or plasticity are direct responses to increased environmental challenge.

The problem is that learning and plasticity tend to be complex adaptations. To build those complex adaptations takes a long time and many changes. That slow process means that one would have to measure the change in evolutionary pressure over the long period during which learning or plasticity evolve. Ultimately, one would need direct measures of environmental pressure in periods of weak challenge and in periods of strong challenge, and measures of the time course of change in learning or plasticity. That is hard to do.

A simpler prediction concerns the rapid increase in the reaction norm in response to an environmental challenge. If a population encounters a novel or extreme environmental challenge, it may respond by broadening its reaction norm. A broader reaction norm can be achieved simply by increasing the tendency for stochastic perturbations during development—a breakdown in the normal homeostatic processes that keep phenotypes within narrow bounds. This prediction about the broadening of the reaction norm follows immediately from Holland's [12] contrast between exploitation and exploration and the quote given in the text by Haldane [11]. Many specific models analyze how the reaction norm responds to environmental challenge. Lande [39] provided new models and a good overview of the literature.

In this case, one simply needs to measure the change in environmental pressure over short periods of time and the associated change in the reaction norm of the population. It would also be interesting to measure how a broader reaction norm acts synergistically with genetic variation to speed the evolutionary response.

VI. CONCLUSION

Evolutionary theory emphasizes genetic variation as the source of evolutionary novelty. By the standard theory, the usual sequence would be a novel environmental challenge, genetic variation either already present or arising *de novo*, and evolutionary response to the novel environment by change in gene frequency.

In this classical evolutionary theory, genetics provides the source of phenotypic variation on which natural selection acts. By contrast, development may generate the

phenotypic novelty that initiates adaptation to environmental challenge. The sequence would be novel environmental challenge, initial survival by those individuals with a phenotypic norm of reaction that overlaps the new fitness surface, and subsequent adaptation by genetic variants from those phenotypes that survive the initial challenge.

West-Eberhard [1] traces the theoretical foundations of this topic from the late 19th century. Since that time, the idea that developmental processes may play a key role in initiating adaptation has never been popular. Evolutionary change is usually tied in thought to genetic change. Nonheritable phenotypic variation by itself is therefore usually believed not to accelerate evolutionary rate.

The original theories of learning, developmental plasticity, and reaction norms have always understood the relations between genotype, phenotype, environment, and evolutionary change. However, the jargon from those theories is thick: the Baldwin effect, genetic assimilation, reaction norms, hopeful monsters, niche construction, and environmentally induced evolution. Each variant theory invoked special environmental conditions, developmental processes, and interactions with genetics. And each in its own way jostled with the ghost of Lamarck. A casual observer could be forgiven for steering clear of the whole mess. Wisdom suggested to wait for clear empirical examples. Induction still dominates mainstream thought in biology.

Many years ago, I read Hinton and Nowlan's [8] article and Maynard Smith's [58] related essay on the Baldwin effect. They focused on the essential theoretical point. Learning smooths the fitness surface, changing evolutionary dynamics in a way that greatly accelerates adaptation to novel or extreme environmental challenges. When one views the whole confusing field in that simple light, one sees that all the complexities of the theories and mechanistic details of phenotypic variability ultimately reduce to the same point. Developmental variation smooths the fitness landscape. A smoothed fitness landscape profoundly alters evolutionary dynamics, particularly in response to novel or extreme environmental challenge.

ACKNOWLEDGMENTS

Marsha Rosner's insight into cellular variability and the evolutionary dynamics of cancer stimulated this work. My research is supported by National Science Foundation grant EF-0822399, National Institute of General Medical Sciences MIDAS Program grant U01-GM-76499, and a grant from the James S. McDonnell Foundation. I completed this article while supported by a Hogge-Baer Visiting Professorship in Cancer Research at the University of Chicago.

-
- [1] M. J. West-Eberhard, *Developmental Plasticity and Evolution* (Oxford University Press, New York, 2003).
 - [2] S. A. Frank, "Natural selection. I. Variable environments and uncertain returns on investment," *Journal of Evolutionary Biology* (2011).
 - [3] C. D. Schlichting and M. Pigliucci, *Phenotypic Evolution: A Reaction Norm Perspective* (Sinauer Associates, Sunderland, Massachusetts, 1998).
 - [4] J. M. Baldwin, "A new factor in evolution," *American Naturalist* **30**, 441–451 (1896).
 - [5] E. Mayr, *The Growth of Biological Thought* (Harvard University Press, Cambridge, Massachusetts, 1982).
 - [6] C. H. Waddington, "Canalization of development and the inheritance of acquired characters," *Nature* **154**, 563–565 (1942).
 - [7] C. H. Waddington, "Genetic assimilation of an acquired character," *Evolution* **7**, 118–126 (1953).
 - [8] G. E. Hinton and S. J. Nowlan, "How learning can guide evolution," *Complex Systems* **1**, 495–502 (1987).
 - [9] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing," *science* **220**, 671–680 (1983).
 - [10] C. J. Geyer and E. A. Thompson, "Annealing markov chain monte carlo with applications to ancestral inference," *Journal of the American Statistical Association* **90**, 909–920 (1995).
 - [11] J. B. S. Haldane, *The Causes of Evolution* (Cornell University Press, Ithaca, NY, 1932).
 - [12] J. H. Holland, *Adaptation in Natural and Artificial Systems* (University of Michigan Press, Ann Arbor, Michigan, 1975).
 - [13] J. A. Coyne, N. H. Barton, and M. Turelli, "A critique of Sewall Wright's shifting balance theory of evolution," *Evolution* **51**, 643–671 (1997).
 - [14] S. Gavrillets, *Fitness Landscapes and the Origin of Species* (Princeton University Press, Princeton, NJ, 2004).
 - [15] L. W. Ance, "Undermining the Baldwin expediting effect: Does phenotypic plasticity accelerate evolution?" *Theoretical Population Biology* **58**, 307–319 (2000).
 - [16] G. P. Wagner and L. Altenberg, "Complex adaptations and the evolution of evolvability," *Evolution* **50**, 967–976 (1996).
 - [17] M. Kirschner and J. Gerhart, "Evolvability," *Proceedings of the National Academy of Sciences USA* **95**, 8420–8427 (1998).
 - [18] M. Pigliucci, "Is evolvability evolvable?" *Nature Reviews Genetics* **9**, 75–82 (2008).
 - [19] E. Rajon and J. Masel, "Evolution of molecular error rates and the consequences for evolvability," *Proceedings of the National Academy of Sciences USA* **108**, 1082–1087 (2011).
 - [20] R. J. Woods, J. E. Barrick, T. F. Cooper, U. Shrestha, M. R. Kauth, and R. E. Lenski, "Second-order selection for evolvability in a large *Escherichia coli* population," *Science* **331**, 1433–1436 (2011).
 - [21] S. Wright, "Evolution in Mendelian populations," *Genetics* **16**, 97–159 (1931).
 - [22] R. Lande, "Genetic variation and phenotypic evolution

- during allopatric speciation,” *American Naturalist* **116**, 463–479 (1980).
- [23] M. C. Whitlock, “Founder effects and peak shifts without genetic drift: adaptive peak shift occur easily when environments fluctuate slightly,” *Evolution* **51**, 1044–1048 (1997).
- [24] G. De Jong, “Quantitative genetics of reaction norms,” *Journal of Evolutionary Biology* **3**, 447–468 (1990).
- [25] S. Gavrillets and S. M. Scheiner, “The genetics of phenotypic plasticity. V. Evolution of reaction norm shape,” *Journal of Evolutionary Biology* **6**, 31–48 (1993).
- [26] R. W. Anderson, “Learning and evolution: a quantitative genetics approach,” *Journal of Theoretical Biology* **175**, 89–101 (1995).
- [27] S. A. Frank, “The design of natural and artificial adaptive systems,” in *Adaptation*, edited by M. R. Rose and G. V. Lauder (Academic Press, San Diego, California, 1996) pp. 451–505.
- [28] P. Turney, “Myths and legends of the Baldwin effect,” in *Proceedings of the Workshop on Evolutionary Computing and Machine Learning, at the 13th International Conference on Machine Learning (ICML-96)* (Bari, Italy, 1996) pp. 135–142.
- [29] G. Mayley, “Landscapes, learning costs, and genetic assimilation,” *Evolutionary Computation* **4**, 213–234 (1997).
- [30] P. Turney, D. Whitley, and R. W. Anderson, “Evolution, learning, and instinct: 100 years of the Baldwin effect,” *Evolutionary Computation* **4**, iv–viii (1997).
- [31] M. Pigliucci, *Phenotypic Plasticity: Beyond Nature and Nurture* (Johns Hopkins University Press, Baltimore, Maryland, 2001).
- [32] S. H. Rice, “A general population genetic theory for the evolution of developmental interactions,” *Proceedings of the National Academy of Sciences USA* **99**, 15518–15523 (2002).
- [33] B. K. Hall, “Baldwin and beyond: organic selection and genetic assimilation,” in *Evolution and Learning: The Baldwin Effect Reconsidered*, edited by B. H. Weber and D. J. Depew (MIT Press, Cambridge, Massachusetts, 2003) pp. 141–167.
- [34] T. D. Price, A. Qvarnström, and D. E. Irwin, “The role of phenotypic plasticity in driving genetic evolution,” *Proceedings of the Royal Society of London. Series B: Biological Sciences* **270**, 1433–1440 (2003).
- [35] R. Mills and R. A. Watson, “On crossing fitness valleys with the Baldwin effect,” in *Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, edited by L. M. Rocha, L. S. Yaeger, M. A. Bedau, D. Floreano, R. L. Goldstone, and A. Vespignani (MIT Press, Cambridge, Massachusetts, 2006) pp. 493–499.
- [36] M. Pigliucci, C. J. Murren, and C. D. Schlichting, “Phenotypic plasticity and evolution by genetic assimilation,” *Journal of Experimental Biology* **209**, 2362–2367 (2006).
- [37] E. Crispo and M. Rausher, “The Baldwin effect and genetic assimilation: revisiting two mechanisms of evolutionary change mediated by phenotypic plasticity,” *Evolution* **61**, 2469–2479 (2007).
- [38] R. Suzuki and T. Arita, “Repeated occurrences of the Baldwin effect can guide evolution on rugged fitness landscapes,” in *Proceedings of the First IEEE Symposium on Artificial Life (IEEE-ALIFE’07)* (IEEE Press, Hoboken, NJ, 2007) pp. 8–14.
- [39] R. Lande, “Adaptation to an extraordinary environment by evolution of phenotypic plasticity and genetic assimilation,” *Journal of Evolutionary Biology* **22**, 1435–1446 (2009).
- [40] L. M. Chevin and R. Lande, “When do adaptive plasticity and genetic evolution prevent extinction of a density-regulated population?” *Evolution* **64**, 1143–1150 (2010).
- [41] L. M. Chevin, R. Lande, and G. M. Mace, “Adaptation, plasticity, and extinction in a changing environment: towards a predictive theory,” *PLoS Biology* **8**, e1000357 (2010).
- [42] S. Gavrillets, “Rapid transition towards the division of labor via evolution of developmental plasticity,” *PLoS Computational Biology* **6**, e1000805 (2010).
- [43] F. Aubret and R. Shine, “Genetic assimilation and the postcolonization erosion of phenotypic plasticity in island tiger snakes,” *Current Biology* **19**, 1932–1936 (2009).
- [44] A. M. Bell and G. E. Robinson, “Behavior and the dynamic genome,” *Science* **332**, 1161–1162 (2011).
- [45] H. Rubin, “The significance of biological heterogeneity,” *Cancer and Metastasis Reviews* **9**, 1–20 (1990).
- [46] I. R. Booth, “Stress and the single cell: intrapopulation diversity is a mechanism to ensure survival upon exposure to stress,” *International Journal of Food Microbiology* **78**, 19–30 (2002).
- [47] E. R. Sumner and S. V. Avery, “Phenotypic heterogeneity: differential stress resistance among individual cells of the yeast *saccharomyces cerevisiae*,” *Microbiology* **148**, 345–351 (2002).
- [48] T. Yomo, Y. Ito, K. Sato, and K. Kaneko, “Phenotypic fluctuation rendered by a single genotype and evolutionary rate,” *Physica A: Statistical Mechanics and its Applications* **350**, 1–5 (2005).
- [49] S. V. Avery, “Microbial cell individuality and the underlying sources of heterogeneity,” *Nature Reviews Microbiology* **4**, 577–587 (2006).
- [50] M. Niepel, S. L. Spencer, and P. K. Sorger, “Non-genetic cell-to-cell variability and the consequences for pharmacology,” *Current Opinion in Chemical Biology* **13**, 556–561 (2009).
- [51] S. L. Spencer, S. Gaudet, J. G. Albeck, J. M. Burke, and P. K. Sorger, “Non-genetic origins of cell-to-cell variability in trail-induced apoptosis,” *Nature* **459**, 428–432 (2009).
- [52] S. J. Altschuler and L. F. Wu, “Cellular heterogeneity: do differences make a difference?” *Cell* **141**, 559–563 (2010).
- [53] K. Kaneko, “Characterization of stem cells and cancer cells on the basis of gene expression profile stability, plasticity, and robustness,” *BioEssays* **33**, 403–413 (2011).
- [54] S. A. Frank, *Immunology and Evolution of Infectious Disease* (Princeton University Press, Princeton, NJ, 2002).
- [55] K. M. Murphy, P. Travers, and M. Walport, *Janeway’s Immunobiology*, 7th ed. (Garland Science, Boca Raton, Florida, 2007).
- [56] E. Loh, J. J. Salk, and L. A. Loeb, “Optimization of DNA polymerase mutation rates during bacterial evolution,” *Proceedings of the National Academy of Sciences USA* **107**, 1154–1159 (2010).
- [57] P. B. Rainey, “Evolutionary genetics: the economics of mutation,” *Current Biology* **9**, R371–R373 (1999).
- [58] J. Maynard Smith, “When learning guides evolution,” *Nature* **329**, 761–762 (1987).